

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION Unclassified			1b RESTRICTIVE MARKINGS N/A	
2a SECURITY CLASSIFICATION AUTHORITY N/A			3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release, distribution unlimited.	
2b DECLASSIFICATION/DOWNGRADING SCHEDULE N/A			5 MONITORING ORGANIZATION REPORT NUMBER(S) N/A	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) N/A AD-A225 816				
6a. NAME OF PERFORMING ORGANIZATION Stanford University, Knowledge Systems Laboratory		6b OFFICE SYMBOL (If applicable)		7a NAME OF MONITORING ORGANIZATION Space and Naval Warfare Systems Command, Department of the Navy
6c ADDRESS (City, State, and ZIP Code) 701 Welch Road Building C Palo Alto, CA 94304		7b ADDRESS (City, State, and ZIP Code) Washington, D.C. 20363-5100		
8a NAME OF FUNDING/SPONSORING ORGANIZATION, Defense Advanced Research Projects Agency		8b OFFICE SYMBOL (If applicable)		9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00039-86-C-0033
8c ADDRESS (City, State, and ZIP Code) 1400 Wilson Boulevard Arlington, VA 22209-2308		10 SOURCE OF FUNDING NUMBERS		
		PROGRAM ELEMENT NO	PROJECT NO	TASK NO
		WORK UNIT ACCESSION NO		
11 TITLE (Include Security Classification) Knowledge-Based Systems Research: Final Report (Unclassified)				
12 PERSONAL AUTHOR(S) Edward Feigenbaum, Paul Rosenbloom, Robert Engelmores				
13a. TYPE OF REPORT Final Technical	13b TIME COVERED FROM 8/22/86 TO 3/31/90		14. DATE OF REPORT (Year, Month, Day) 8/24/90	15 PAGE COUNT 17
16 SUPPLEMENTARY NOTES				
17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	Artificial Intelligence, Blackboard Systems, Constraint Satisfaction, Machine Learning, Knowledge Acquisition, Simulation, Logic, Design	
19 ABSTRACT (Continue on reverse if necessary and identify by block number) Knowledge-Based Systems Research covers research on the following topics: Blackboard model of problem solving; constraint satisfaction; machine learning and knowledge acquisition; symbolic simulation; logic-based systems with self-awareness; SOAR, an architecture for general intelligence and learning; software design by redesign.				
20 DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL John Machado			22b. TELEPHONE (Include Area Code) (202) 692-3966	22c. OFFICE SYMBOL

DTIC
ELECTE
AUG 28 1990
S B D

Final Report

Knowledge-Based Systems Research

Edward A. Feigenbaum
Robert S. Engelmores
Paul C. Rosenbloom

Sponsored by
Space and Naval Warfare Systems Command
Computational Systems and Technology Division
(SPAWAR 3212)
under Contract No. N00039-86C-0033
Effective Date of Contract: August 22, 1986
Contract Expiration Date: March 31, 1990

Report Date: August 24, 1990
Contractor: Knowledge Systems Laboratory, Stanford University
Principal Investigator: Edward A. Feigenbaum
415-723-4878

The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of the Navy or the U. S. Government.

Knowledge-Based Systems Research Final Report

This is the final report on Knowledge-Based Systems Research, under the terms of Contract # N00039-86-C-0033, issued by the Space and Naval Warfare Systems Command. The official start date of the contract was August 22, 1986, although research actually began on October 10, 1985, when pre-award costs were approved. The contract ended on March 31, 1990.

Knowledge-Based Systems Research covers research on the following topics:

1. Blackboard model of problem solving
2. Constraint satisfaction
3. Machine learning and knowledge acquisition
4. Symbolic simulation
5. Logic-based systems with self-awareness
6. SOAR, an architecture for general intelligence and learning
7. Software design by redesign

Task Objectives

Blackboard Systems: Deepen scientific understanding and improve the power and range of application of this framework for problem-solving. Design and implement an environment in which blackboard systems and applications can easily be built, providing facilities for controlling reasoning, explaining actions, and learning new problem solving strategies.

Constraint Satisfaction: Produce tools for constructing symbolic constraint satisfaction programs, and analyze their performance to evaluate strengths and limitations in problem-solving.

Machine Learning and Knowledge Acquisition: Develop robust machine learning programs to be integrated in a variety of intelligent systems. Determine criteria for application of machine learning techniques to various problem-solving architectures.

Symbolic Simulation: Develop concepts and software for symbolic simulation of complex, dynamic systems. Build representation schemes for time course of events and develop animation techniques to improve user understanding of modeled behavior.

Logic-based systems with self-awareness: Explore and implement new reasoning techniques, viz. uncertain reasoning for resolution, theory formation based on measures of probability and simplicity, efficiency-enhancing theory reformulation, and counterfactual implication.

SOAR: Formulate a general abstraction-planning mechanism and a general problem-space creation mechanism. Demonstrate SOAR on several large tasks. Perform experiments to determine the sufficiency of chunking as a general learning mechanism. Complete revision of SOAR into a robust, efficient, user-friendly system, capable of helping the user give it new tasks.

Software design by redesign: Develop a design environment to support and complement the skills of human software designers. Develop a representation for software designs and specifications with multiple perspectives. Build a prototype library of software modules. Develop methods for modification and propagation of design changes. Develop a user interface language and protocol.

For	
[]	
[]	
[]	
ed	
tion	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Technical Problems

No significant technical problems were encountered.

General Methodology

The methodology employed in all phases of this research has been engineering-oriented and experimental. That is, research problems are explored by designing, building and experimenting with programs that serve to test underlying theories.

Technical Results

In this final report we give a brief synopsis of research in each of the topic areas. An extensive bibliography of publications that were written under this contract is presented in the appendix.

1. Blackboard model of problem solving

A blackboard-based problem solving architecture, BB1, was completed and released (version v2.0). A more general framework, called BB*, was also completed and released (version v1.0). Both systems are available in CommonLisp. Copies of each implementation were distributed to several academic and industrial research organizations that had been awaiting them. BB1's system-building support and run-time capabilities were extended in several areas: (a) declarative representation of large bodies of factual and heuristic knowledge; (b) integration of multiple reasoning skills in a single system; (c) dynamic control under real-time constraints. We have implemented these application-independent components: (a) declarative representation for device structure, function, faults, and repairs; (b) reasoning modules for associative and model-based diagnosis; and (c) an asynchronous communications interface. These components were demonstrated in the GUARDIAN system for managing medical-intensive-care patients in real time (see below).

BB* permits the integration of BB1 with application specific frameworks which include knowledge and terminology about specific tasks. One such framework is ACCORD, which tailors the system for the task of arranging objects to satisfy constraints. The utility of ACCORD has been demonstrated for system design, implementation, and performance in two application systems: the PROTEAN system for protein structure modeling and the SIGHTPLAN system for designing construction-site layouts. ACCORD (version v1.0), written in CommonLisp, has also been released.

A new application, called GUARDIAN, was developed for managing patients in medical intensive care. This application permits exploring such issues as interfacing a reasoning system to the external environment, modeling dynamic systems, and using real-time constraints on reasoning and problem solving.

GUARDIAN Demo-2 was completed. It performs the following functions in real time: (a) monitors and filters data from a respirator according to dynamically specified I/O filters; (b) dynamically adjusts I/O filters so that the system monitors data at the maximum rate, given its current reasoning activities; (c) classifies data in meaningful temporal episodes; (d) probabilistically diagnoses abnormal episodes; (e) uses generic flow models to explain the respiratory system structure, function, and faults underlying particular diagnoses; and (f) uses generic flow models to hypothesize alternative diagnoses. In developing Demo 2, we addressed research issues related to selective attention, representation and use of generic structure/function

models, incremental diagnosis over time, integrating multiple reasoning activities, and coordination with a dynamic external environment.

A two-part paper, *Blackboard Systems*, by H. Penny Nii, was published in the *AI MAGAZINE*, Volumes 7, No. 2 and Volume 7, No. 3. The paper explains what is meant by the blackboard model of problem solving, defines a computational framework for blackboard systems, traces the intellectual history of the "blackboard" ideas, and describes and contrasts half-a-dozen blackboard application programs.

A full-length book, also entitled *Blackboard Systems*, was completed and published by Addison-Wesley in 1988. The book, edited by R. Englemore and A. Morgan, is a collection of most of the major technical papers on blackboard systems through 1986.

2. Constraint satisfaction

Constraint satisfaction techniques were explored in the context of two different application domains, budgetary constraints for financial resource management and spatial constraints for molecular structure determination.

For the financial resource management application we built a system called FRM for creating proposal budgets and matching records of expenditures with budgets. In FRM, all financial management domain knowledge is represented in the constraint form. The preferences, policies, and procedures that the FRM system uses in resource management are written in the form of symbolic constraints that are applied to appropriate items in a developing budget. The constraint management program propagates constraints and attempts to satisfy them sequentially.

For the spatial constraint application, we developed a system called PROTEAN for determining the three-dimensional structure of protein molecules from NMR data. PROTEAN incorporates high speed geometric calculations for spatial arrangement problems and uses BB1 knowledge sources to control the direction and order of constraint propagation. PROTEAN has been used successfully to partially assemble various proteins subject to experimental and theoretical constraints. The system was later expanded with the addition of new functionality to refine a protein at the atomic level of resolution. To demonstrate and validate these capabilities, constraints generated from a known protein (cytochrome-b562) were used to reconstruct the structure at the atomic level of detail, yielding a family of atomic structures as a reasonable depiction of the molecular structure.

3. Machine learning and knowledge acquisition

Investigations were carried out in a number of different sub-areas of machine learning and knowledge acquisition, including learning by induction, apprenticeship learning, intelligent editors, text learning, communication patterns between knowledge engineers and experts, and learning by analogy.

A number of experiments were conducted with a rule induction program, RL, which showed how intelligent selection of instances increases the efficiency of learning from examples. Other experiments with RL were conducted in the domain of high energy physics to show the generality of RL's method. These experiments demonstrated that a simulation program can effectively provide training examples for a learning program. Another series of experiments were performed with RL to determine the effectiveness of learning from examples of different types, especially for the concept of a "near miss."

A prototype program that learns by watching how an expert solves a diagnosis problem was implemented. Designed Meta-OPAL, a framework for building intelligent editors and began implementation. The REFEREE program which critiques articles on clinical trials, was rewritten from EMYCIN into the object-oriented language KEE to allow interactive acquisition of new knowledge from published journal articles. A long-term study of communication patterns in the interaction of experts and knowledge engineers was initiated. Research was completed, under M. Genesereth's direction, on a deductive approach to analogical reasoning.

4. Symbolic simulation

We explored qualitative simulation as a means of representing theories of gene regulation in molecular biology, and developed several new qualitative representation and reasoning techniques. This work is fully documented in the Ph.D. dissertation of Peter Karp, entitled "Hypothesis Formation and Qualitative Reasoning in Molecular Biology."

During the summer we made progress in developing a process-oriented simulation system for the tryptophan operon. The thesis describes, among other things, a process-oriented simulation system for the tryptophan operon. The important objects in this biological system and the important interactions between them were identified, thus enabling the simulation of all important processes in the operon such as chemical binding events, and protein translation. In addition to the library of biological objects, the system includes a means of describing initial experimental conditions, and a partial library of processes which describe interactions between objects. An interpreter was implemented which executes these processes in a fashion similar to a production system. The interpreter can be run to compute the topology of a reaction network for a given set of initial conditions. Construction of the knowledge base raised a number of important issues in constructing large knowledge bases.

The use of simulation techniques for automated discovery of process mechanisms was investigated. A process-oriented simulation system was developed for generating and testing hypotheses about biological regulation mechanisms.

A successful prototype of an expert system that combines rule-based symbolic reasoning with numerical optimization and simulation, for diagnostic troubleshooting of a particle beam accelerator, was also completed.

5. Logic-based systems with self-awareness

Work on advanced inference techniques stressed formalization of the notion of irrelevance, the investigation of the expressive and computational properties of defeasible inference, and the application of counterfactual inference to planning.

Work on multi-valued logics focussed on the implementation of a backward chainer. A multi-valued inference engine was developed that has been specialized to yield a first-order theorem prover, an assumption-based truth maintenance system, and a default reasoner (both with and without justification information). The program has been used to successfully solve a variety of problems in the simulation and diagnosis of digital hardware

Work on advanced reasoning progressed on four fronts: control of reasoning, reasoning about action, theory of abstraction reformulations for computational efficiency, and the implementation and integration of these methods into the MRS system.

This research includes work on 1) automatic compilation of MRS axioms into efficient code; 2) implementation of automatic control in MRS; 3) work on theory and implementation of abstraction reformulations; and 4) work on reasoning about derived information over time .

6. SOAR, an architecture for general intelligence and learning

The Soar project is focused on the task of developing, applying, and evaluating an architecture usable as the basis of an integrated intelligent agent. By its nature, such an effort spans a large part of AI. The segment of the Soar project that has been funded through Darpa contract N00039-86C-0033 to the Stanford Knowledge Systems Laboratory has focused on making contributions in the areas of abstraction, integrated models of learning, self-compiling expert systems, and production system technology. In addition to being contributions in their own right, each of these pieces contributes to the overall capabilities of the Soar system. These four areas of contribution are summarized below, followed by a few words on the integration of rules and case-based/analogical reasoning.

In the area of abstraction, we have been examining how abstraction can arise automatically, and how it can integrate synergistically with problem solving (or planning) and learning. We have developed a family of methods that can automatically simplify task descriptions, learn from the simplified task descriptions, and transfer what is learned to reduce the effort required to accomplish the full task. In the process, these methods increase the tractability of the learning method, and enable the acquisition of generalized rules. The family of methods includes a very weak method of broad applicability, plus stronger iterative and assumption-counting methods. These methods have been implemented in Soar, and shown to be effective in three disparate domains: a puzzle domain (eight puzzle), a planning domain (Strips-like mobile robot), and an expert system domain (computer configuration). Three of the key ways that these methods differ from earlier abstraction methods are: (1) that they automatically determine what to abstract (and when to abstract it); (2) that they dynamically and incrementally abstract the behavior of the system, rather than statically abstracting the structure of the system (during a preprocessing phase); and (3) that they, because of their implementation in Soar, lead to the acquisition of abstract rules -- which, in combination, form an abstract plan -- that can transfer to later tasks.

In the area of integrated models of learning, we have been attempting to understand the span of learning behaviors possible via chunking in Soar. Our hypothesis is that this covers the full range of learning behaviors required of an intelligent agent. Towards this end we have developed and implemented a framework which integrates together three of the principal forms of machine learning: rote learning, induction from multiple examples, and explanation-based learning. The framework shows how all three can arise from chunking, if the variety of searches over which chunking works is sufficiently wide. Also, as part of this integration, we have shown how chunking can acquire three of the key types of knowledge required of an intelligent system: procedural knowledge, episodic knowledge, and declarative knowledge. This work has provided novel solutions to several of the key problems in explanation-based learning (EBL), such as how EBL can learn truly new things, how EBL can be integrated with other forms of learning, and how to decide what concepts EBL should be used to learn about.

In the area of self-compiling expert systems, we have been investigating how to construct "first principles" expert systems in Soar, and how they can then automatically compile themselves, via chunking (EBL), into shallower, more efficient systems. Our early work -- prior to this contract -- was in the area of computer configuration -- a classical construction expert system. The work on abstraction in computer configuration (described above) was an attempt to reduce the searches that are required before such learning can occur, and to further speed up the compilation by learning more general rules (so fewer overall need be learned). Our other major effort was an

investigation into a very different style of expert system -- a classical classification expert system. Neomycin-Soar is a medical diagnosis expert system that is a reimplementations of the Neomycin system (which is itself a reimplementations of the Mycin system). It is the largest expert system to have been built in Soar, at approximately 5000 rules. The results of this effort were: (1) a demonstration that both classification and construction tasks map cleanly onto Soar's basic generate-select-execute cycle; (2) a demonstration that a Neomycin-like "first principles" approach could be automatically compiled down to a tight question-asking loop, where learned rules are used directly to determine what question to ask next, and what the consequences of the answer are; and (3) that a major source of the learning-induced speedup comes from rules that avoid dead-ends in the reasoning process.

In the area of production system technology, we have been working towards efficient and bounded production matchers. The production matcher is in the inner loop of Soar (and other systems), and both efficiency and boundedness are required for effective performance, especially as we head towards real-time domains. On the efficiency side, we have developed an optimized RETE-based matcher for Soar that, among other things, incorporates a sophisticated condition reorderer. We have also implemented a TREAT-based matcher for Soar, and have shown experimentally that RETE outperforms TREAT under a set of identifiable assumptions that are met by Soar. On the boundedness side, we have developed a space of novel production system formulations that admit polynomial-time match algorithms. The unique-attribute formulation, which has a linear-time match algorithm, has been shown to eliminate the problem of expensive chunks -- where the acquisition of new rules degrades performance because of their exponential match cost. While none of the new formulations yet appears ideal, several have proved to be interesting contenders for replacing the current Ops5-style formulation. They also reveal an interesting connection between production matching and constraint satisfaction problems.

One effort that was spawned by supported work on analogy in Soar, and that continued to be supported under this contract though it was continued outside of the Soar system, was on the integration of rules with case-based/analogical reasoning. This work has so far made a number of contributions, most notably in the area of how to use cases that do not come predigested with rationales for their various actions. A rational-reconstruction technique has been developed whereby plausible explanations are generated for the cases (from weak background theories), and then used as the basis for analogical transfer to new situations. A second contribution is the development of a metric that determines when to accept the results of the available rules, or when to instead use analogy to override what the rules suggest.

7. Software design by redesign

During the course of our research in blackboard systems, a collaborative research project was initiated between KSL and Kestrel Institute to investigate the design of software architectures. Software designers often design new systems by remembering and modifying a "similar" conceptual design they used in the past. The objective of this project is to develop a system that provides aid in this process.

During the period of this contract we developed an initial knowledge representation scheme for describing software architectures, and built a prototype knowledge base (library) of blackboard architectures and their components. A scheme for representing functional modules from various perspectives was proposed. Using this scheme generic components of blackboard architectures were built up from the perspectives of functional decomposition hierarchy, control flow, and data flow. Some "redesign operators" were identified that are intended to transform the components in the library into ones with different characteristics that would fit into a developing architecture (that is, instantiate the library module for a specific architecture).

A prototype graphic interface for the design/redesign environment was specified. The specification includes browsing through the knowledge base from different viewpoints, and constructing a SADT-like diagram from the information in the knowledge base. The SADT diagram is expected to become one method the user can use to describe system requirements and changes to system design. SADT's expressibility was modified and extended to enable the display of the contents of the knowledge base.

Important Findings and Conclusions

The research performed under this contract resulted in a large number of evolutionary improvements in our understanding of the power and limitations of knowledge-based systems. These findings are contained in an extensive library of technical reports, which are listed in the appendix.

Significant Hardware Development

N/A

Special Comments

The Knowledge Systems Laboratory has a long tradition of attention to technology transfer, and we continued to emphasize that activity during the course of this contract. In each topic area the research results have been documented in books, journal articles, technical reports and tutorial notes. Where appropriate, the knowledge-based technology has been transferred to other universities, industrial research laboratories and defense companies in the form of software. An example of that transfer process took place with Boeing Computer Services. Computer scientists from BCS spent time at the KSL learning to use the blackboard system BB1. They then installed the software on their equipment at BCS and developed an extended version, called ERASMUS. This system was then used by personnel at the Boeing Military Air Center to develop a prototype control system for a head-up display on the Advanced Tactical Fighter.

Implications for Further Research

This research, which continues previous research on knowledge-based systems, has demonstrated the power of knowledge as the performance driver in intelligent systems. The "knowledge is power" hypothesis can be elevated to a principle, based on the evidence accumulated over the past 15 years. However, we can now see more clearly the limitations of current generation knowledge-based systems – their over-specialization and their brittleness when faced with novel problems – and the implications therein for further research. If future systems are to be robust problem solvers like the human experts they emulate, they must be given far greater and more basic knowledge than they now have. We need to find ways to represent and use knowledge about the world at multiple levels of detail, to be able to formulate qualitative as well as quantitative models of objects or processes and use these models to predict and explain behavior when compiled knowledge doesn't apply to the situation at hand. Further research in building large, shared, multi-use knowledge bases is clearly implied.

Appendix: Bibliography

BLACKBOARD MODEL OF PROBLEM SOLVING

Daube, F. and Hayes-Roth, B. 1988. FIRST: A Case-Based Redesign System in the BB1 Blackboard Architecture. In Proceedings of the AAAI-88 Workshop on Case-Based Reasoning. Also Technical Report, KSL 88-32, August 1988, 7 pages.

Engelmore, R. and Morgan, A., Eds. 1988. *Blackboard Systems*. Wokingham England: Addison-Wesley Ltd.

Garvey, A., Hewett, M., Johnson, M. V., Schulman, R., Hayes-Roth, B. 1986. BB1 User Manual - Interlisp Version, Technical Report, KSL 86-60, October 1986, 68 pages

Garvey, A., Hewett, M., Johnson, M. V., Schulman, R., Hayes-Roth, B. 1986. BB1 User Manual - Common Lisp Version, Technical Report, KSL 86-61, October 1986, 72 pages.

Garvey, A. and Hayes-Roth, B. 1987. Implementing Diverse Forms of Control Knowledge in Multiple Controls, Technical Report, KSL 87-40, June 1987, 32 pages.

Garvey, A., Cornelius, C., and Hayes-Roth, B. 1987. Computational Costs versus Benefits of Control Reasoning, Technical Report, KSL 87-11, June 1987, 6 pages.

Garvey, A. and Hayes-Roth, B. 1989. An Empirical Analysis of Explicit vs. Implicit Control Architectures. In *Blackboard Architectures and Applications*, eds. V. Jagannathan, R. Dodhiawala and L. S. Baum, 43-56. San Diego: Academic. Also Technical Report, KSL 88-21, March 1988, 20 pages.

Hayes-Roth, B., Johnson, M. V., Garvey, A., and Hewett, M. 1986. A Layered Environment for Reasoning about Action. *Journal of Artificial Intelligence in Engineering, Special Issue on Blackboard Systems*, October 1986. Also Technical Report, KSL 86-38, or STAN-CS-87-1147, April 1987, 63 pages

Hayes-Roth, B., Johnson, M. V., Garvey, A., and Hewett, M. 1986. Applications of BB1 to Arrangement-Assembly Tasks. *Journal of Artificial Intelligence in Engineering, Special Issue on Blackboard Systems*, October 1986.

Hayes-Roth, B. 1987. A Multi-Processor Interrupt-Driven Architecture for Adaptive Intelligent Systems, Technical Report, KSL 87-31, June 1987, 19 pages.

Hayes-Roth, B. 1987. Dynamic Control Planning in Adaptive Intelligent Systems, Technical Report, KSL 87-67, November 1987, 7 pages.

Hayes-Roth, B. 1988. Making Intelligent Systems Adaptive. In *Architectures for Intelligence*, eds. A. Newell and K. VanLehn. Hillsdale, N.J.: Lawrence Erlbaum. Also Technical Report, KSL 88-64, September 1988, 24 pages.

Hayes-Roth, B., Hewett, R., and Seiver, A. 1988. Diagnostic Explanation using Generic Models, Technical Report, KSL 88-20, March 1988, 10 pages.

Hayes-Roth, B., Hewett, M., Johnson, M. V., and Garvey, A. 1988. ACCORD: A Framework for a Class of Design Tasks, Technical Report, KSL 88-19, March 1988, 12 pages.

- Hayes-Roth, B. 1989. Control in Blackboard Systems. In *Blackboard Architectures and Applications*, eds. V. Jagannathan, R. Dodhiawala and L. S. Baum, 3-7. San Diego: Academic.
- Hayes-Roth, B., Hewett, M., Washington, R., Hewett, R., and Seiver, A. 1989. Distributing Intelligence Within a Single Individual. In *Distributed Artificial Intelligence, Volume 2*, eds. L. Gasser and M. N. Huhns, 385-412. San Mateo: Morgan Kaufmann.
- Hayes-Roth, B., Washington, R., Hewett, R., Hewett, M., and Seiver, A. 1989. Intelligent Real-Time Monitoring and Control. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 243-249. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence. Also Technical Report, KSL 89-05, 15 pages.
- Hewett, M. and Hayes-Roth, B. 1987. The BB1 Architecture: A Software Engineering View, Technical Report, KSL 87-10, February 1987, 17 pages.
- Hewett, M. and Hayes-Roth, B. 1989. Real-Time I/O in Knowledge-Based Systems. In *Blackboard Architectures and Applications*, eds. V. Jagannathan, R. Dodhiawala and L. S. Baum, 269-283. San Diego: Academic. Also Technical Report, KSL 88-22, November 1988, 13 pages.
- Hewett, R. and Hayes-Roth, B., and Seiver, A. 1989. Representing Generic Knowledge of Physical Systems for Use in Multiple Contexts. In Proceedings of the AAAI Symposium on Knowledge System Development Tools and Languages, Stanford Univ., March 1989.
- Hewett, R. and Hayes-Roth, B. 1990. Representing and Reasoning About Physical Systems Using Generic Models. In *Formal Aspects of Semantic Networks*, eds. J. Sowa, S. Shapiro and R. Brachman. San Mateo, Calif.: Morgan Kaufmann. Also Technical Report, KSL 90-40, May 1990, 19 pages.
- Johnson, M. V. and Hayes-Roth, B. 1986. Integrating Diverse Reasoning Methods in the BB1 Blackboard Control Architecture, Technical Report, KSL 86-76, December 1986, 17 pages.
- Johnson, M. V. and Hayes-Roth, B. 1988. Learning to Solve Problems by Analogy, Technical Report, KSL 88-01, March 1988, 15 pages.
- Nii, H. P. 1986. Blackboard Systems. *AI Magazine* 7(2):38-53 and 7(3):82-106. Also Technical Report, KSL 86-18, or STAN-CS-86-1123, 86 pages.
- Schulman, R. and Hayes-Roth, B. 1987. ExAct: A Module for Explaining Actions, Technical Report, KSL 87-08, January 1987, 15 pages.
- Schulman, R. and Hayes-Roth, B. 1988. Plan-Based Construction of Strategic Explanations, Technical Report, KSL 88-23, March 1988, 13 pages.
- Schulman, R., and Hayes-Roth, B. 1988. A Generic Approach to Strategic Explanation. In Proceedings of the AAAI Workshop on Explanation, St. Paul, Minnesota, August 1988.
- Tommelein, I. D., Levitt, R. E., and Hayes-Roth, B. 1987. Using Expert Systems for the Layout of Temporary Facilities on Construction Sites. Presented at the Chartered Institute of Building W-65 Symposium, Organization and Management of Construction, Birkshire, U.K., September, 1987.
- Tommelein, I. D., Johnson, M. V., Hayes-Roth, B., and Levitt, R. E. 1987. SIGHTPLAN-A Blackboard Expert System for the Layout of Temporary Facilities on a Construction Site. Presented at the International Federation for Information Processing Working Conference on Expert Systems in Computer-Aided Design, WG5.2, Sydney, Australia, February 16-20, 1987.

Washington, R. and Hayes-Roth, B. 1989. Input Data Management in Real-Time AI Systems. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 250-253. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

CONSTRAINT SATISFACTION

Altman, R. and Buchanan, B. G. 1987. Partial Compilation of Strategic Knowledge. In Proceedings of the Sixth National Conference on Artificial Intelligence, 399-404. Menlo Park, Calif.: American Association for Artificial Intelligence.

Altman, R., Duncan, B., Brinkley, J., Buchanan, B. G., and Jardetzky, O. 1988. Determination of the Spatial Distribution of Protein Structure Using Solution Data. In *NMR Spectroscopy in Drug Research*, eds. J. W. Jaroszewski, K. Schaumberg and H. Kofod, 592-599. Copenhagen: Munksgaard.

Altman, S. 1988. Representing and Editing Constraints: A Case Study in Financial Resource Management, Technical Report, KSL 88-45, June 1988, 42 pages.

Brinkley, J. F., Buchanan, B. G., Altman, R. B., Duncan, B. S., and Cornelius, C. W. 1987. A Heuristic Refinement Method for Spatial Constraint Satisfaction Problems, Technical Report, KSL 87-05, or STAN-CS-87-1142, January 1987, 15 pages.

Brinkley, J., Cornelius, C., Altman, R., Hayes-Roth, B., Lichtarge, O., Duncan, B., Buchanan, B., and Jardetzky, O. 1988. Application of Constraint Satisfaction Techniques to the Determination of Protein Tertiary Structure, Technical Report, KSL 86-28, March 1986, 14 pages.

Brugge, J. A. and Buchanan, B. G. 1987. Evolution of a Knowledge-Based System for Determining Structural Components of Proteins, Technical Report, KSL 87-62, October 1987, 26 pages.

Finger, J. J. 1987. Exploiting Constraints in Design Synthesis, Ph.D. thesis, Dept. of Computer Science, Stanford Univ., March 1987.

Gelman, A., Altman, S., Pallakoff, M., Doshi, K., Manago, C., Rindfleisch, T. C., and Buchanan, B. G. 1988. FRM: An Intelligent Assistant for Financial Resource Management. In Proceedings of the Seventh National Conference on Artificial Intelligence, 31-36. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 88-29, April 1988, 25 pages.

Hayes-Roth, B., Buchanan, B. G., Lichtarge, O., Hewett, M., Altman, R., Brinkley, J., Cornelius, C., Duncan, B., and Jardetzky, O. 1988. PROTEAN: Deriving Protein Structure from Constraints. In *Blackboard Systems*, eds. R. Englemore and A. Morgan, 417-431. Wokingham, England: Addison-Wesley. Also in Proceedings of the Fifth National Conference on Artificial Intelligence, 904-909. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 86-51, March 1986, 21 pages.

Henager, D. E. 1986. Window-Driven Object-Oriented Calculator, Technical Report, KSL 86-64 (Working Paper), March 1986, 56 pages.

Jardetzky, O., Lane, A., Lefevre, J. F., Lichtarge, O., Hayes-Roth, B., Altman, R., and Buchanan, B. G. 1986. A New Method for the Determination of Protein Structures in Solution from NMR. In Proceedings XXIII Congress Ampere, eds. B. Maraviglia, F. De Luca and R. Campanella, 64-69, Rome, Italy. Also Technical Report, KSL 86-42, May 1986, 6 pages.

Washington, R. and Hayes-Roth, B. 1989. Input Data Management in Real-Time AI Systems. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 250-255. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

CONSTRAINT SATISFACTION

Altman, R. and Buchanan, B. G. 1987. Partial Compilation of Strategic Knowledge. In Proceedings of the Sixth National Conference on Artificial Intelligence, 399-404. Menlo Park, Calif.: American Association for Artificial Intelligence.

Altman, R., Duncan, B., Brinkley, J., Buchanan, B. G., and Jardetzky, O. 1988. Determination of the Spatial Distribution of Protein Structure Using Solution Data. In *NMR Spectroscopy in Drug Research*, eds. J. W. Jaroszewski, K. Schaumberg and H. Kofod, 592-599. Copenhagen: Monksgaard.

Altman, S. 1988. Representing and Editing Constraints: A Case Study in Financial Resource Management, Technical Report, KSL 88-45, June 1988, 42 pages.

Brinkley, J. F., Buchanan, B. G., Altman, R. B., Duncan, B. S., and Cornelius, C. W. 1987. A Heuristic Refinement Method for Spatial Constraint Satisfaction Problems, Technical Report, KSL 87-05, or STAN-CS-87-1142, January 1987, 15 pages.

Brinkley, J., Cornelius, C., Altman, R., Hayes-Roth, B., Lichtarge, O., Duncan, B., Buchanan, B., and Jardetzky, O. 1988. Application of Constraint Satisfaction Techniques to the Determination of Protein Tertiary Structure, Technical Report, KSL 86-28, March 1986, 14 pages.

Brugge, J. A. and Buchanan, B. G. 1987. Evolution of a Knowledge-Based System for Determining Structural Components of Proteins, Technical Report, KSL 87-62, October 1987, 26 pages.

Finger, J.J. 1987. Exploiting Constraints in Design Synthesis, Ph.D. thesis, Dept. of Computer Science, Stanford Univ., March 1987.

Gelman, A., Altman, S., Pallakoff, M., Doshi, K., Manago, C., Rindfleisch, T. C., and Buchanan, B. G. 1988. FRM: An Intelligent Assistant for Financial Resource Management. In Proceedings of the Seventh National Conference on Artificial Intelligence, 31-36. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 88-29, April 1988, 25 pages.

Hayes-Roth, B., Buchanan, B. G., Lichtarge, O., Hewett, M., Altman, R., Brinkley, J., Cornelius, C., Duncan, B., and Jardetzky, O. 1988. PROTEAN: Deriving Protein Structure from Constraints. In *Blackboard Systems*, eds. R. Englemore and A. Morgan, 417-431. Wokingham, England: Addison-Wesley. Also in Proceedings of the Fifth National Conference on Artificial Intelligence, 904-909. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 86-51, March 1986, 21 pages.

Henager, D. E. 1986. Window-Driven Object-Oriented Calculator, Technical Report, KSL 86-64 (Working Paper), March 1986, 56 pages.

Jardetzky, O., Lane, A., Lefevre, J. F., Lichtarge, O., Hayes-Roth, B., Altman, R., and Buchanan, B. G. 1986. A New Method for the Determination of Protein Structures in Solution from NMR. In Proceedings XXIII Congress Ampere, eds. B. Maraviglia, F. De Luca and R. Campanella, 64-69, Rome, Italy. Also Technical Report, KSL 86-42, May 1986, 6 pages.

MACHINE LEARNING and KNOWLEDGE ACQUISITION

- Buchanan, B. G., Sullivan, J., Cheng, T.-P., and Clearwater, S. 1988. Simulation-Assisted Inductive Learning. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 552-557. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Cheng, T.-P. 1989. Towards a Formalization of Partial Domain Theories, Landmark Value Selection and Incremental Learning for Inductive Learning Systems, Technical Report, KSL 89-53 (Working Paper), June 1989, 51 pages.
- Clancey, W. J. 1987. Knowledge Engineering Methodology: An Annotated Bibliography of NEOMYCIN Research, Technical Report, KSL 87-58, September 1987, 10 pages.
- Clancey, W. J. 1987. The Knowledge Engineer as Student: Metacognitive Bases for Asking Good Questions. In *Learning Issues for Intelligent Tutoring Systems*, eds. H. Mandl and A. Lesgold. New York: Springer-Verlag. Also Technical Report, KSL 87-12, or STAN-CS-87-1183, January 1987, 30 pages.
- Clancey, W. J. 1987. Diagnosis, Teaching, and Learning: An Overview of GUIDON2 Research, Technical Report, KSL 87-32 (Working Paper), April 1987, 12 pages.
- Fu, L. M. 1985. Learning Object-Level and Meta-Level Knowledge in Expert Systems, Ph.D. diss., Technical Report, KSL 85-44, November 1985, 229 pages.
- Heckerman, D. and Jimison, H. 1989. A Perspective on Confidence and Its Use in Focusing Attention during Knowledge Acquisition. In *Uncertainty in Artificial Intelligence 3*, eds. L. Kanal, T. Levitt and J. Lemmer, 123-131. New York: North-Holland. Also Technical Report, KSL 87-45 (Journal Memo), July 1987, 10 pages.
- Keller, R. M. 1988. Defining Operationality for Explanation-Based Learning. *Artificial Intelligence* 35 (2):227-241. Also Technical Report, KSL 87-60, October 1987, 19 pages.
- Keller, R. M. 1988. Operationality and Generality in Explanation-Based Learning: Separate Dimensions or Opposite Endpoints? In *Proceedings of the AAAI Symposium on Explanation-Based Learning*, 153-157, March 1988, Stanford, CA. Also Technical Report, KSL 88-11, February 1988, 5 pages.
- Keller, R. M. 1988. Learning Approximate Concept Descriptions, Technical Report, KSL 88-57, July 1988, 15 pages.
- Musen, M. A. 1988. Conceptual Models of Interactive Knowledge-Acquisition Tools, Technical Report, KSL 88-16, March 1988, 16 pages.
- Prieditis, A., Dietterich, T. G., Hirsh, H., Kedar-Cabelli, S., Kempiminski, R., Minton, S., and Subramanian, D. 1986. AAAI-86 Learning Papers: Developments and Summaries. *Machine Learning* 2(1): 83-96.
- Rodolitz, N. S. 1987. Guidon Manage - Tutoring for Strategic Knowledge, Technical Report, KSL 87-38, June 1987, 43 pages.
- Subramanian, D. and Smith, D. E. 1988. Knowledge Level Learning: An Alternative View. In *Proceedings of the AAAI Spring Symposium on Explanation-Based Learning*, 196-200, March 1988, Stanford, CA.
- Sullivan, J. 1988. RL3: An Approach to Incremental Rule Learning, Technical Report, KSL 88-52 (Working Paper), June 1988, 24 pages.

Wilkins, D. C. and Buchanan, B. G. 1986. On Debugging Rule Sets When Reasoning Under Uncertainty. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 448-454. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 86-30, or STAN-CS-87-1181, 19 pages.

Wilkins, D. C., Buchanan, B. G., and Clancey, W. J. 1986. Knowledge Base Debugging Using Apprenticeship Learning Techniques. In *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop*, November 1986. Also Technical Report, KSL 86-63, October 1986, 15 pages.

Wilkins, D. C., Clancey, W. J., and Buchanan, B. G. 1987. Knowledge Base Refinement Using Abstract Control Knowledge, Technical Report, KSL 87-01, or STAN-CS-87-1182, January 1987, 9 pages.

Wilkins, D. C., Clancey, W. J., and Buchanan, B. G. 1987. Using and Evaluating Differential Modeling in Intelligent Tutoring and Apprentice Learning Systems. In *Intelligent Tutoring Systems: Lessons Learned*, 257-278. Hillsdale, N. J.: Lawrence Erlbaum. Also Technical Report, KSL 86-62, or STAN-CS-87-1175, January 1987, 37 pages.

SYMBOLIC SIMULATION

Karp, P. D. 1988. A Process-Oriented Model of Bacterial Gene Regulation, Technical Report, KSL 88-18, June 1988, 14 pages.

Karp, P. and Friedland, P. 1989. Coordinating the Use Of Qualitative and Quantitative Knowledge in Declarative Device Modeling. In *Artificial Intelligence, Modeling and Simulation*, eds. L. E. Widman, D. H. Helman and K. Loparo, 189-206. New York: John Wiley. Also Technical Report, KSL 87-09, January 1988, 22 pages.

Karp, P. D. 1989. Hypothesis Formation and Qualitative Reasoning in Molecular Biology (thesis), Technical Report, KSL 89-52, June 1989, 337 pages.

Round, A. D. 1987. QSOPS: A Workbench Environment For The Qualitative Simulation of Physical Processes, Technical Report, KSL 87-37, May 1987, 5 pages. To appear in: *Proceedings of the European Simulation Multiconference*.

Round, A. D. 1989. Knowledge-Based Simulation. In *The Handbook of Artificial Intelligence, Volume IV*, eds. A. Barr, P. R. Cohen and E. A. Feigenbaum, 415-518. Reading, Mass.: Addison-Wesley.

LOGIC-BASED SYSTEMS WITH SELF-AWARENESS

Ginsberg, M. L. 1986. Counterfactuals. *Artificial Intelligence* 30 (1): 35-80.

Ginsberg, M. L. 1986. Implementing Probabilistic Reasoning. In *Uncertainty in Artificial Intelligence*, eds. L. N. Kanal and J. F. Lemmer, 331-338. Amsterdam: North-Holland.

Ginsberg, M. L. 1986. Bilattices. Presented at the Second International Conference on Lattice Theory and Universal Algebra. Also submitted to *Journal of Symbolic Logic*. Unpublished.

Ginsberg, M. L. 1986. Multi-Valued Logics. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 243-247. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 86-29, April 1986, 13 pages.

Ginsberg, M. L. 1986. Possible Worlds Planning. Presented at the 1986 Workshop on Planning and Reasoning About Action. Also Technical Report, KSL 86-37 (Working Paper), April 1986, 13 pages.

Ginsberg, M. L. 1987. Steps Toward a General-Purpose Inference Engine. Unpublished.

Ginsberg, M. L. and Smith, D. E. 1987. Reasoning About Action I: A Possible Worlds Approach. *Artificial Intelligence* 35 (2):165-195. Also Technical Report, KSL 86-65, May 1987, 25 pages. Short version appeared in Proceedings of the Sixth National Conference on Artificial Intelligence, 212-217. Menlo Park, Calif.: American Association for Artificial Intelligence.

Ginsberg, M. L. and Smith, D. E. 1987. Reasoning About Action II: The Qualification Problem. *Artificial Intelligence* 35 (3):311-342. Also Technical Report, KSL 86-66, May 1987, 28 pages. Short version appeared in Proceedings of the Sixth National Conference on Artificial Intelligence, 212-217. Menlo Park, Calif.: American Association for Artificial Intelligence.

Smith, D. E. 1987. Controlling Backward Inference. *Artificial Intelligence* 39(2):145-208. Also Technical Report, KSL 86-68, March 1987, 67 pages.

Subramanian, D. and Feigenbaum, J. 1986. Factorization in Experiment Generation. In Proceedings of the Fifth National Conference on Artificial Intelligence, 518-522. Menlo Park, Calif.: American Association for Artificial Intelligence.

Subramanian, D. and Genesereth, M. R. 1986. Reformulation. In Proceedings of the AAAI/SIGART Workshop on Knowledge Compilation, Oregon, September 1986.

Subramanian, D. and Genesereth, M. R. 1987. The Relevance of Irrelevance. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 416-422. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Subramanian, D. 1989. A Theory of Justified Reformulations, Ph.D. thesis, Technical Report, STAN-CS-90-1260, March 1989. Also a summary of it appears in *Change of Representation and Inductive Bias*, ed. P. Benjamin, 147-168. Hingham, Mass.: Kluwer Academic, 1990.

SOAR, AN ARCHITECTURE FOR GENERAL INTELLIGENCE AND LEARNING

Golding, A., Rosenbloom, P., and Laird, J. E. 1987. Learning General Search Control from Outside Guidance. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 334-337. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Golding, A. and Rosenbloom, P. S. 1989. Combining Analytical and Similarity-Based CBR. In Proceedings of the Case-Based Reasoning Workshop, Pensacola, FL.

Nayak, P., Gupta, A., and Rosenbloom, P. 1988. Comparison of the Rete and Treat Production Matchers for SOAR (A Summary). In Proceedings of the Seventh National Conference on Artificial Intelligence, 693-698. Menlo Park, Calif.: American Association for Artificial Intelligence.

Rosenbloom, P. S. and Newell, A. 1986. The Chunking of Goal Hierarchies: A Generalized Model of Practice. In *Machine Learning: An Artificial Intelligence Approach, Volume II*, eds. R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, 247-288. Los Altos, Calif.: Morgan Kaufmann.

Rosenbloom, P. S., Laird, J. E., and Newell, A. 1986. *Universal Subgoalting and Chunking: The Automatic Generation and Learning of Goal Hierarchies*. Hingham, Mass.: Kluwer Academic.

- Rosenbloom, P. S., Laird, J. E., and Newell, A. 1987. Knowledge Level Learning in SOAR. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, 499-504. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Rosenbloom, P. S. 1989. A Symbolic Goal-Oriented Perspective on Connectionism and SOAR. In *Connectionism in Perspective*, eds. R. Pfeifer, Z. Schreter, F. Fogelman-Soulie, and L. Steels, 245-263. Amsterdam: Elsevier (North-Holland).
- Rosenbloom, P. S., Laird, J. E., and Newell, A. 1989. The Chunking of Skill and Knowledge. In *Working Models of Human Perception*, eds. B. A. G. Elsendoorn and H. Bouma, 391-410. London: Academic Press.
- Rosenbloom, P. S. and Aasman, J. 1990. Knowledge Level and Inductive Uses of Chunking (EBL). In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 821-827. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Rosenbloom, P. S., Newell, A. and Laird, J. E. 1990. Towards the Knowledge Level in SOAR: The Role of the Architecture in the Use of Knowledge. In *Architectures for Intelligence*, ed. K. VanLehn. Hillsdale, N.J.: Erlbaum. In press.
- Rosenbloom, P. S., Laird, J. E., Newell, A. and McCarl, R. 1991. A Preliminary Analysis of the SOAR Architecture as a Basis for General Intelligence. *Artificial Intelligence*. To appear.
- Scales, D. J. 1986. Efficient Matching Algorithms for the SOAR/OPS5 Production System, Technical Report, KSL 86-47, or STAN-CS-86-1124, June 1986, 50 pages
- Tambe, M. and Rosenbloom, P. S. 1989. Eliminating Expensive Chunks by Restricting Expressiveness. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 731-737. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Tambe, M. and Rosenbloom, P. S. 1990. A Framework for Investigating Production System Formulations with Polynomially Bounded Match. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 693-700. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Tambe, M. and Rosenbloom, P. S. 1990. Production System Formulations with Non-Combinatoric Match: Some Results and Analysis. Unpublished.
- Tambe, M., Newell, A. and Rosenbloom, P. S. 1990. The Problem of Expensive Chunks and Its Solution by Restricting Expressiveness. *Machine Learning*. To appear.
- Unruh, A., Rosenbloom, P. S., and Laird, J. E. 1987. Dynamic Abstraction Problem Solving in SOAR. In *Proceedings of the AOG/AAAIC Joint Conference*, 245-256, Dayton, OH.
- Unruh, A. and Rosenbloom, P. 1989. Abstraction in Problem Solving and Learning. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 681-687. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Washington, R. and Rosenbloom, P. S. 1989. Applying Problem Solving and Learning to Diagnosis. Unpublished.

SOFTWARE DESIGN BY REDESIGN

- Peyton, L. 1988. A Transformational Approach to Software Redesign, KSL 88-37 (Working Paper), May 1988, 23 pages.

REASONING WITH UNCERTAINTY

Grosof, B. N. 1986. An Inequality Paradigm for Probabilistic Knowledge. In *Uncertainty in Artificial Intelligence*, eds. L. N. Kanal and J. F. Lemmer, 259-275. Amsterdam: North-Holland.

Grosof, B. N. 1986. Evidential Confirmation as Transformed Probability. In *Uncertainty in Artificial Intelligence*, eds. L. N. Kanal and J. F. Lemmer, 153-166. Amsterdam: North-Holland.

Grosof, B. N. 1986. Non-Monotonicity in Probabilistic Reasoning. In *Proceedings of the Second AAAI/RCA Workshop on Uncertainty in Artificial Intelligence*, 91-98. Menlo Park, Calif.: American Association for Artificial Intelligence.

Heckerman, D. E. and Horvitz, E. J. 1987. On the Expressiveness of Rule-Based Systems for Reasoning with Uncertainty. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, Vol.1, 121-126. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 87-53 (Journal Memo), August 1987, 7 pages.

Horvitz, E. J. 1987 Reasoning About Beliefs and Actions Under Computational Resource Constraints. In *Proceedings of the Third Workshop on Uncertainty in Artificial Intelligence*, AAAI, Seattle, Washington, March 1989. Also Technical Report, KSL 87-29, (Journal Memo), March 1987, 26 pages.

Horvitz, E. J. 1988. Reasoning Under Varying and Uncertain Resource Constraints. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 111-116. Menlo Park, Calif.: American Association for Artificial Intelligence. Also Technical Report, KSL 88-35, April 1988, 6 pages.

Wellman, M. P. and Heckerman, D. E. 1987. The Role of Calculi in Uncertain Reasoning. In *Proceedings of the Third Workshop on Uncertainty in Artificial Intelligence*, 355-362. Mountain View, Calif.: Association for Uncertainty in Artificial Intelligence. Also Technical Report, KSL 87-46 (Journal Memo), July 1987, 12 pages.

GENERAL

Barsalou, T. 1987. An Object-Based Interface to a Relational Database System, Technical Report, KSL 87-41, June 1987, 10 pages.

Buchanan, B. G. 1988. Artificial Intelligence As An Experimental Science. In *Aspects of Artificial Intelligence*, ed. J. H. Fetzer. Boston: Kluwer Academic. Also Technical Report, KSL 87-03, January 1987, 41 pages.

Buchanan, B. G. and Smith, R. G. 1988. Fundamentals of Expert Systems, Technical Report, KSL 88-09 (Working Paper), March 1988, 33 pages.

Buchanan, B. G. and Shortliffe, E. H. 1989. Advances in Expert Systems (A White Paper and Commentary). Presented to DARPA, February, 1989. Also Technical Report, KSL 89-33 (Working Paper), April 1989, 21 pages.

Buchanan, B. G., Bobrow, D., Davis, R., McDermott, J., and Shortliffe, E. H. 1989. Research Directions in Knowledge-Based Systems. *Annual Review of Computer Science 4*, ed. J. Traub. Also Technical Report, KSL 89-71, October 1989, 15 pages.

Feigenbaum, E. A., McCorduck, P., and Nii, H. P. 1988. *The Rise of the Expert Company*. New York: Times Books.

Horvitz, E. J. 1987. Inference under Varying Resource Limitations, Technical Report, KSL 87-16 (Journal Memo), February 1987, 16 pages.

Horvitz, E. J. 1987. Problem-Solving Design: Reasoning about Computational Value, Tradeoffs, and Resources. In Proceedings of the National Aeronautics And Space Administration Artificial Intelligence Forum, Mountain View, CA. Technical Report, KSL 87-64 (Journal Memo), November 1987, 19 pages.

Karp, P. D. and Wilkins, D. C. 1989. An Analysis of the Distinction Between Deep and Shallow Expert Systems. *International Journal of Expert Systems* 1(1). Also Technical Report, KSL 89-10, February 1989, 36 pages.

Lederberg, J. 1987. How DENDRAL Was Conceived and Born. In Proceedings of the ACM Symposium on the History of Medical Informatics, National Library of Medicine, November 1987. Also Technical Report, KSL 87-54, August 1987, 19 pages.

Selig, L. J. 1987. An Expert System Using Numerical Simulation and Optimization To Find Particle Beam Line Errors, Technical Report, KSL 87-36, May 1987, 39 pages.

Walker, M. G. 1986. How Feasible is Automated Discovery? *IEEE Expert* 2(1):70-82. Also Technical Report, KSL 86-35, April 1986, 24 pages.

Zygmunt, A. 1988. SOLACE: Systems Optimization Laboratory's Automated Computational Expertise, Technical Report, KSL 88-39, May 1988, 55 pages.